

CLAIMS

- 1 1. A method for constructing an extended array of storage devices that enables storage of
 2 more than one parity block per stripe on a single device of the array, the method com-
 3 prising the steps of:
 4 dividing each device into blocks;
 5 organizing the blocks into stripes across the devices, wherein each stripe contains
 6 a first number f_1 of parity blocks from each of a second number f_2 of parity storage de-
 7 vices of the extended array; and
 8 storing data blocks on a third number n of data storage devices of a first set of de-
 9 vices that is disjoint from a second set of devices storing the parity blocks, to thereby
 10 construct the extended array adapted to enable storage of more than one parity block per
 11 stripe on a single device, wherein f_1 times f_2 equals n , and wherein the extended array is
 12 recoverable from any one or two storage device failures.
- 1 2. The method of Claim 1 further comprising the step of assigning the data blocks to
 2 parity sets of the array, wherein each data block is assigned to two parity sets and
 3 wherein each parity set contains at least one data block and a parity block.
- 1 3. The method of Claim 2 further comprising the step of defining the parity sets, includ-
 2 ing the data and parity blocks, within stripes in the first and second sets of devices.
- 1 4. The method of Claim 3 wherein the first set of devices is a set of data disks and the
 2 second set of devices is a set of parity disks, and wherein each data disk contains one or
 3 more blocks of a stripe and each parity disk contains the same or fewer blocks from the
 4 stripe.
- 1 5. The method of Claim 4 further comprising the step of numbering each parity set from
 2 0 to $n-1$, where n is the number of data disks.

1 6. The method of Claim 5 wherein the step of assigning comprises the step of precluding
 2 assignment of certain parity deltas to the data blocks, where a parity delta is a difference
 3 in numerical value, modulo n , of the two parity sets to which a data block is assigned.

1 7. The method of Claim 6 wherein the step of storing comprises the step of storing the
 2 parity blocks separated by the precluded parity deltas on the single device of the second
 3 set.

1 8. The method of Claim 7 further comprising the step of configuring the extended array
 2 with a predetermined number of devices.

1 9. The method of Claim 8 wherein the predetermined number of devices is an even num-
 2 ber n and wherein one of the precluded parity deltas is $n/2$.

1 10. The method of Claim 1 wherein the step of storing comprises the step of creating an
 2 unbalanced stripe array of data and parity blocks across the devices.

1 11. The method of Claim 10 further comprising the step of combining a plurality of un-
 2 balanced stripe arrays to thereby produce a single balanced super-stripe array across the
 3 devices.

1 12. A system for constructing an extended storage array that enables storage of more
 2 than one parity block per stripe on a storage device of the array, the system comprising:
 3 a plurality of storage devices, each device divided into blocks that are further or-
 4 ganized into stripes, wherein each stripe contains a first number f_1 of parity blocks from
 5 each of a second number f_2 of parity storage devices of the extended array;
 6 a storage operating system including a device storage layer configured to imple-
 7 ment a concentrated parity technique that stores data blocks on a third number n of data
 8 storage devices of a first set of devices that is disjoint from a second set of devices stor-

9 ing the parity blocks, wherein f_1 times f_2 equals n , and that groups the stripes of data
 10 blocks into a contiguous set of stripes from which the parity blocks of a corresponding
 11 group of stripes of parity blocks is computed to thereby construct the extended array
 12 adapted to enable storage of more than one parity block within a group of parity blocks
 13 on a single device; and

14 a processing element configured to execute the operating system to thereby in-
 15 voke storage access operations to and from the array in accordance with the concentrated
 16 parity technique.

1 13. The system of Claim 12 wherein the device storage layer assigns the data blocks to
 2 parity sets of the array, each parity set containing at least one data block and a parity
 3 block, while precluding assignment of certain parity deltas to the data blocks.

1 14. The system of Claim 13 wherein the device storage layer further stores the parity
 2 blocks separated by the precluded parity deltas on a single device of the second set.

1 15. The system of Claim 14 wherein the storage devices are disks and wherein the device
 2 storage layer is a RAID layer.

1 16. The system of Claim 15 wherein the RAID layer is implemented in logic circuitry.

1 17. The system of Claim 14 wherein the storage devices are one of video tape, magnetic
 2 tape, optical, DVD and bubble memory devices.

1 18. The system of Claim 14 wherein the storage devices are media adapted to store in-
 2 formation contained within the data and parity blocks.

1 19. Apparatus for constructing an extended array of disks that enables storage of more
 2 than one parity block per stripe on a single disk of the array, the apparatus comprising:
 3 means for dividing each disk into blocks;

means for organizing the blocks into stripes across the disks, wherein each stripe contains a first number f_1 of parity blocks from each of a second number f_2 of parity disks of the extended array; and

means for storing data blocks on a third number n of data disks of a first set of disks that is disjoint from a second set of disks storing the parity blocks to thereby construct the extended array adapted to enable storage of more than one parity block per stripe on a single disk, wherein f_1 times f_2 equals n , and wherein the extended array is recoverable from any one or two storage device failures.

20. The apparatus of Claim 19 wherein the means for storing comprises means for creating an unbalanced stripe array of data and parity blocks across the disks.

21. The apparatus of Claim 20 further comprising means for combining a plurality of unbalanced stripe arrays to thereby produce a single balanced super-stripe array across the disks.

22. A computer readable medium containing executable program instructions for constructing an extended array of disks that enables storage of more than one parity block per stripe on a single disk of the array, the executable program instructions comprising program instructions for:

configuring the extended array with a predetermined number of disks;

dividing each disk into blocks;

organizing the blocks into stripes across the disks, wherein each stripe contains a first number f_1 of parity blocks from each of a second number f_2 of parity disks of the extended array and wherein each disk contains one or more blocks from each stripe;

assigning the data blocks to parity sets of the array, wherein each parity set contains at least one data block and a parity block; and

storing data blocks on a third number n of data disks of a first set of disks that is disjoint from a second set of disks storing the parity blocks, to thereby construct the extended array adapted to enable storage of more than one parity block per stripe on a sin-

15 gle disk, wherein f_1 times f_2 equals n and wherein the extended array is recoverable from
16 any one or two storage device failures.

1 23. The computer readable medium of Claim 22 wherein the program instruction for as-
2 signing comprises the program instruction for precluding assignment of certain parity
3 deltas to the data blocks.